

Cooling Solutions & Considerations for High Performance Computers

White Paper 6

Revision 1

By Rich Whitmore

Executive summary

The recent advances in high performance computing have driven innovations not only in the computer hardware market but also in several supporting industries such as storage, power and cooling. All of these technologies play an important support role in the HPC industry and perhaps none has received more focus recently than cooling.

Introduction

As the uses for High Performance Computing (HPC) continue to expand, the requirements for operating these advanced systems has also grown. The drive for higher efficiency data centers is a subject closely tied to a building's Power Usage Effectiveness (PUE) which is defined as: $(\text{Total Facility Energy}) / (\text{IT Equipment Energy})$. HPC clusters and high density compute racks are consuming power densities of up to 100 kW per rack, sometimes higher, with an estimated average density of 35 kW per rack. Building owners, Co-location facilities, Enterprise data centers, Web Scale companies, Governments, Universities and National Research Laboratories are struggling to upgrade cooling infrastructure to not only remove the heat generated by these new computer systems but also to reduce or eliminate their effects on building energy footprints or PUE. The new and rapid adaptation of "Big Data" HPC systems in industries such as oil and gas research, financial trading institutions, web marketing and others is further highlighting the need for efficient cooling due to the fact that the majority of the world's computer rooms and data centers are not equipped nor prepared to handle the heat loads generated by current and next generation HPC systems. If one considers that 100% of the power consumed by an HPC system is converted to heat energy, it's easy to see why the removal of this heat in an effective and efficient manner has become a focal point of the industry.

Investigating Cooling Solutions for High Performance Computer Systems

Submersion Cooling

New high performance computer chips have the capability of allowing an HPC system designer to develop special clusters that can reach 100 kW per rack and exceed almost all available server cooling methods currently available. Submersion cooling systems offer baths or troughs filled with a specially designed nonconductive dielectric fluid allowing entire servers to be submerged in the fluid without risk of electrical conductance across the computer circuits. These highly efficient systems can remove up to 100% of the heat generated by the HPC system. The



Submersion Cooling Bath

heat, once transferred to the dielectric fluid can then be easily removed via a heat exchanger, pump and closed loop cooling system.

To be applied, the traditional data center is typically renovated to accept the new submersion cooling system. Legacy cooling equipment such as CRAC's, raised flooring and vertical server racks are replaced by the submersion baths and updated closed loop warm water cooling system. These baths lay horizontally on the floor which provides a new vantage point for IT personnel although at the cost of valuable square footage. Servers are modified either by the owner or a 3rd party by removing components that would be negatively affected by the dielectric fluid such as hard drives and other components that might not be warranted by original equipment manufacturers (OEMs). Special consideration should be paid to future server refresh options considering that such a monumental shift in infrastructure greatly limits OEM server options in the future and limits the overall uses for a server room with dedicated submersion cooling technology only.


While submersion cooling offers extreme efficiencies for the world's most extreme HPC systems, the general scarcity of such HPC systems combined with required infrastructure upgrades and maintenance challenges poses an issue for market wide acceptance at this time.



The worldwide High Performance Computing (HPC) market is expected to grow at an 8.3% CAGR reaching \$44 billion in 2020. The global HPC market will generate \$220 Billion in revenues over the period 2015-2020.

Direct to Chip, On-Chip Cooling

Direct to Chip or On-Chip cooling technology has made significant advances recently in the HPC industry. Small heat sinks are attached directly to the computer CPU's and GPU's creating high efficiency close coupled server cooling. Up to 70% of the heat from the servers is collected by the Direct to Chip heat sinks and transferred through a system of small capillary tubes to a Coolant Distribution Unit (CDU). The CDU then transfers the heat to a separate closed loop cooling system to reject the heat from the computer room. The balance of the heat, 30% or more is rejected to the existing room cooling infrastructure.



The warm water cooling systems commonly used for Direct to Chip cooling are generally considered cooling systems that don't utilize a refrigeration plant such as closed loop dry coolers (similar to large radiators) and cooling towers and have been recently quantified by the American Society of Heating Refrigeration and Air-Conditioning Engineers (ASHRAE) to produce "W-3 or W-4" water temperatures or water ranging from 2°C - 46°C (36°F-115°F). These systems draw significantly less energy than a typical refrigerated chiller system and provide adequate heat removal for direct to chip cooling systems as they can operate with cooling water supply temperatures in the W3 - W4 range.

Direct to chip cooling solutions can also be used to reclaim low grade water heat that if repurposed and used correctly can improve overall building efficiencies and PUE. The advantages of this form of heat recovery is limited by the ability of the building's HVAC system to accept it. HVAC building design varies around the world. Many parts of Europe can benefit from low grade heat recovery because of the popular use of water based terminal units in most buildings. In contrast most of the North American HVAC building designs use central forced air heating and cooling systems with electric reheat terminal boxes, leaving little use for low grade heat recovery from a direct to chip or on chip cooling system. The feasibility to distribute reclaimed warm water should also be studied in conjunction with building hydronic infrastructure prior to use.

A recent study performed by the Ernest Orlando Lawrence Berkeley National Laboratory titled "Direct Liquid Cooling for Electronic Equipment" concluded that the best cooler performance achieved by a market leading direct to chip cooling system reached 70% under optimized laboratory conditions. This leaves an interesting and possibly counterproductive result for such systems because large amounts of heat from the

computer systems must still be rejected to the surrounding room which must then be cooled by more traditional, less efficient means such as Computer Room Air Conditioners (CRAC) and or Computer Room Air Handlers (CRAH) (Diagram A). To understand better the new result of deploying a direct or on chip cooling system, one must consider the HPC cluster as part of the overall building energy footprint which can then be directly tied to the building PUE. Considering a 35 kW HPC rack with direct to chip cooling will reject at least 10.5 kW (30%) of heat to the computer room and an average HPC cluster consists of 6 racks of compute racks (excluding high density storage arrays), a direct to chip and or on chip cooling system will reject at least 60 kW of heat load to a given space. Utilizing the most common method of rejecting this residual heat by CRAC or CRAH results in a significant setback in original efficiency gains.

Additional challenges are presented when considering the actual infrastructure needed inside the data center and more importantly inside the server rack itself when considering an on-chip cooling system. In order to bring warm water cooling to the chip level, water must be piped inside the rack through a large number of small hoses which in turn feed the small direct to chip heat exchangers/pumps. While small by nature, the scale of which these are installed, leaves an IT staff looking at the back of a rack filled with large numbers of hoses as well as a distribution header for connecting to the inlet and outlet water of the cooling system.

Direct to chip cooling systems are tied directly to the mother boards of the HPC cluster and are designed to be more or less permanent. Average HPC clusters are refreshed (replaced) every 3-5 years typically based on demand or budget. With that in mind, one must consider the costs of replacing the direct to chip or on chip cooling infrastructure along with the server refresh.

Direct to chip cooling offers significant advances in efficiently cooling today's high performance computer clusters, however once put into the context of a larger computer room or building, one must consider the net result of overall building performance, cost implications and useful life on total ROI.

Active Rear Door Heat Exchangers

Active Rear Door Heat Exchangers (ARDH) have grown in popularity among those manufacturing and using HPC clusters and high density server racks. An ARDH's ability to remove 100% of the heat from the server rack with little infrastructure change offers advantages in system efficiency and convenience. These systems are commonly rack agnostic and replace the rear door of any industry standard server rack. They utilize an array of high efficiency fans and tempered cooling water to remove heat from the computer system. The Electronically Commutated (EC) fans work to match server air flow rate in CFM to ensure all heat is removed from the servers. An ARDH uses clean water or a glycol mix between 57F-75F which is often readily available in most data centers and if not, can be produced by a chilled water plant, closed loop cooling systems such as cooling towers, dry fluid coolers, or a combination of these systems. Utilizing an ARDH allows for high density server racks to be installed in existing computer rooms such as Co-location Facilities or legacy data centers with virtually no infrastructure changes and zero effect on surrounding computer racks.



Active Rear Door Heat Exchangers are capable of removing up to 75 kW per compute rack, offering users large amounts of scale as clusters go through multiple refresh cycles. These systems once deployed also offer advantages to owners by monitoring internal server rack temperatures and external room temperatures, ensuring that a heat neutral environment is maintained.

Recent laboratory tests by server manufacturers have found that the addition of an ARDH actually reduces fan power consumption of the computers inside that rack, more than offsetting the minimal power consumption of the ARDH fan array. While counterintuitive at first glance, in depth study has shown that ARDH fans assist the server's fans allowing them to draw less energy and perform better even at high density workloads. Tests also show that hardware performance improves resulting in increased server life expectancy.

ARDH's offer full access to the rear of the rack and can be installed in both top and bottom water feed configuration offering further flexibility to integrate into new or existing facilities with or without the use of raised floors.

Conclusion

As densities in high performance computer systems continue to increase, the applied method of cooling them has become increasingly more important. Continued expansion for the use of HPC and high density servers in ever broadening applications will see these types of computers installed in more traditional data centers bringing overall data center PUE conversations into focus where in the past, HPC clusters have been largely overlooked or exempt from the conversation due to their previously more limited use in government labs and large research facilities. There are several reliable technologies currently available to cool today's HPC and high density server systems and one must select an efficient and practical system that works within the associated building's cooling infrastructure, future refresh strategy and budget. The engineering and design plans for cooling these advanced computers should take place prior to or in parallel with purchasing of the computer system as the cooling system itself is now often leveraged to ensure optimal and guaranteed computer performance.

About the author

Rich Whitmore is the President and CEO of Motivair Corporation, manufacturer of mission critical chillers and the ChilledDoor® rack cooling system. A graduate in Mechanical Engineering from the Rochester Institute of Technology, his career has focused on cooling some of the world's most advanced data centers and computers as well as pioneering advances in high efficiency, close coupled server cooling systems.

References

- ASHRAE TC 9.9, Thermal Guidelines for Data Processing Environments, 3rd Edition
- Direct Liquid Cooling for Electronic Equipment by Henry Coles and Steve Greenberg, Ernest Orlando Lawrence Berkley National Laboratory
- Worldwide High Performance Computing (HPC) Market Forecast 2015-2020, Market Research Media 2014

Contact us

For feedback and comments about the content of this white paper:

info@motivaircorp.com

For information regarding the ChilledDoor Rack Cooling System:

www.chilleddoor.com

For information regarding Motivair Corporation:

www.motivaircorp.com

©2015 Motivair Corporation. All rights reserved.

